

# The Publication Choice Problem

Haichuan Wang<sup>1</sup>, Yifan Wu<sup>2</sup>, Haifeng Xu<sup>3</sup>

<sup>1</sup>Harvard University

<sup>2</sup>Microsoft Research

<sup>3</sup>University of Chicago

haichuan\_wang@g.harvard.edu, yifan.wu2357@gmail.com, haifengxu@uchicago.edu

## Abstract

Researchers strategically choose where to submit their work in order to maximize its impact, and these publication decisions in turn determine venues' impact factors. To analyze how individual publication choices both respond to and shape venue impact, we introduce a game-theoretic framework - coined the *Publication Choice Problem* - that captures this two-way interplay. We show the existence of a pure-strategy equilibrium in the Publication Choice Problem and its uniqueness under binary researcher types. Our characterizations of the equilibrium properties offer insights about what publication behaviors better indicate a researcher's impact level.

Through equilibrium analysis, we further investigate how labeling papers with "spotlight" affects the impact factor of venues in the research community. Our analysis shows that competitive venue labeling top papers with "spotlight" may decrease the overall impact of other venues in the community, while less competitive venues with "spotlight" labeling have an opposite impact.

**Full version of this paper** —

<https://arxiv.org/abs/2511.13678>

**Dataset** — <https://github.com/Haichuan23/Pub-choice-data>

## 1 Introduction

How to choose publication venues is a strategic choice of researchers because they derive rewards from publications—particularly in prestigious venues—which are central to a research career. As a result, researchers are often rational about their publication strategies in response to venue impacts. In turn, the average impacts of publication venues are also subject to the strategic behaviors of researchers and can co-evolve over time with such behaviors. For instance, the rising popularity of machine learning conferences has been associated with a perceived decline in the *average* impact of publications. Consequently, researchers might instead choose smaller conferences that are considered more selective and more impactful.

Inspired by the feedback loop in the job market signaling game (Spence 1978), we propose the Publication Choice Problem and analyze its dynamics and equilibrium. The

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

game consists of a continuum of researchers (agents) and a set of publication venues, defined as follows.

- **Researchers (agents).** Each researcher in our model is a principal investigator (PI), parameterized by a type representing impact level and endowed with a uniform time budget. The type captures the researcher's productivity, research taste, etc., which we assume is determined by a researcher's past actions and does not vary over the time considered in our model. The uniform time budget assumption reflects that PIs generally have a similar amount of time available for research. Students help convert the PI's time into publications. A large group of students does not increase the PI's time budget but reduces her publication costs.
- **Publication venues.** Each venue has a venue impact and an intrinsic competitiveness level that reflects on the cost of a publication for each researcher type on the venue. The venue impact is defined as the weighted average of researcher types who publish papers on the venue. The publication cost captures the publication venue's acceptance rate, location, preference over topics, etc., and does not vary over the time considered in our model.

Researchers strategize to gain utility from publishing in high-impact venues. The utility of the researcher can result from the recognition of her paper by a venue where high-impact type researchers publish. Each researcher solves a utility maximization problem, with the actions being the number of publications in each venue, and subject to the constraint of total time budget. For example, consider an Publication Choice Problem with two venues, venue 1 and 2. The researcher observes impacts of the two venues as  $v_1 = 0.2$  and  $v_2 = 0.7$ . Assuming the researcher uses up her time budget, she can choose to publish either (A) 3 papers on venue 1 and 1 paper on venue 2, or (B) 1 paper on venue 1 and 2 papers on venue 2. The utility gained from strategy (A) is lower than that of strategy (B), leading the researcher to prefer the latter.

Upon observing venues' impacts from the last round, the researchers modify their publication strategies. In each round in our model, researchers simultaneously choose the number of papers to publish in each venue upon observing the impact factor of all the venues. After all researchers publish their papers, the impact of each venue is updated to

	# publications		Utility
	venue 1	venue 2	
Impact	0.2	0.7	
Action 1	3	1	$3 \cdot 0.2 + 0.7 = 1.3$
Action 2	1	2	$2 \cdot 0.7 + 0.2 = 1.6$

Table 1: The utility of two actions. Each action is a vector of the number of publications (# publications) on each venue.

the average type of researchers who publish in that venue. These updates are then observed by the researchers in the next round. An equilibrium is reached if venue impacts do not change over rounds, i.e., when venue impacts align with researcher actions.

By modeling and analyzing the Publication Choice Problem, we are able to reveal the effect of researchers’ strategic publication behaviors on the publication venues’ impact. We summarize our main results as below.

**Equilibrium Existence** A pure-strategy equilibrium always exists (Proposition 3.3). Moreover, under the binary-type setting (researchers are either high-impact or low-impact), there exists a unique pure-strategy equilibrium (Theorem 4.1).

**Indicators of researcher’s impact level** The total number of publications is not monotone in researcher’s impact (Proposition 3.4). However, the number of publications on the best venue is monotone in researcher’s impact (Theorem 3.1).

**Spotlight Labeling** When a venue labels some papers as “spotlight” papers, there exists a threshold effect on the venue impact (Theorem 4.3): a less competitive venue improves the impact of all venues by setting up spotlight labeling, while a more competitive venue decreases the impact of all venues. In fact, the spotlight papers divert the impact in a research community from regular venues to the special spotlight papers, thus reducing overall impacts.

The rest of the paper proceeds as follows. In Section 2, we introduce the model of Publication Choice Problem. In Section 3, we characterize equilibrium properties. Section 3.1 analyzes the best-response dynamic. Section 3.2 introduces our key assumption on publication costs, provides justification for it, and offers sanity checks of the model. Section 3.3 shows the equilibrium properties, including the existence of an equilibrium and the signaling effect of publication numbers in researcher types. We focus on the binary-type setting in Section 4. We are able to show the equilibrium uniquely exists in Section 4.1. With the equilibrium uniqueness, in Section 4.2, we focus on the venue organizer’s spotlight design problem and examine how switching to spotlight labeling may change the equilibrium outcome. In Section 5, we summarize our contribution and discuss future work.

## 1.1 Related Work

**The Science of Science.** Our paper contributes to the “science of science” literature, which studies the science of research publication such as quantifying research impacts

(Wang, Song, and Barabási 2013; Frank et al. 2019; Perry and Reny 2016), designing better review systems (Su 2021; Zhang, Schoenebeck, and Su 2024; Stelmakh, Shah, and Singh 2021a; Lipton and Steinhardt 2019; Aziz, Micha, and Shah 2023; Boehmer, Brederbeck, and Nichterlein 2022; Fromm et al. 2021; Payan and Zick 2022) and understanding incentives/competition in academia (Zhang et al. 2022; Heckman and Moktan 2020; Ductor et al. 2020; Manzoor and Shah 2021; Stelmakh, Shah, and Singh 2021b). There exists an extensive literature that models strategic behavior in research publication and attempts to understand and improve peer review processes (Zhang et al. 2022; Shah 2021; Wu et al. 2024; Zhang, Schoenebeck, and Su 2024; Stelmakh, Shah, and Singh 2021a; Lipton and Steinhardt 2019; Wright, Thornton, and Leyton-Brown 2015; Jecmen et al. 2020; Meir et al. 2021). Our work also studies strategic behavior, but our focus is mainly on the impact on the research community with multiple publication venues. Along this line, the most relevant to us is the concurrent work of Ductor et al. (2020), which similarly studies the evolution of publication venues in a research community. However, the goal of Ductor et al. (2020) is completely different from ours, hence their detailed modeling. Motivated by publication conventions in economics, Ductor et al. (2020) study the evolution of a research community with a single general-purpose publication venue and several field venues. However, motivated by the publication patterns in machine learning (or AI, generally), our work does not focus specifically on the impact of a general-purpose venue<sup>1</sup>. Instead, our model considers the impact tiers of different venues, which applies to publication patterns in a general research field, such as sub-fields in computer science and economics. Despite similarities in some model components (e.g., the two-side model structure and venue quality modeling), on modeling choice, Ductor et al. (2020) models discrete submission strategy of general-purpose or field venues, with multiple equilibria arising from this discrete strategy space. Our work models the submission numbers to different venues as a vector in a continuous submission space. We are thus able to characterize the closed-form best response. Besides, the data sets and technical results are different.

**Large Population Games.** Our Publication Choice Problem is a game with a large population, which has been studied in multiple research fields including mean field games (Lasry and Lions 2007; Aumann 1975; Lauriere et al. 2022; Perrin et al. 2020) and (non-atomic) congestion games (Milchtaich 1996; Friedman 1996; Blonski 1999; Roughgarden and Tardos 2004). Conceptually, of particular relevance is the congestion games where too many players picking a certain action will render that action bad due to congestion. Our Publication Choice Problem bears some similarity but differs fundamentally in at least two key aspects, which renders techniques in congestion games (e.g. potential function) inapplicable to our problem. First, the utility from each venue does not depend on the total number of players but rather on their average impact. Second and more

<sup>1</sup>ML conferences include diverse topics, making research fitness less important for AI researchers.

importantly, Publication Choice Problem allows heterogeneous player types while congestion game does not.

## 2 A Model of Publication Choice

**Conventional Notations.** Throughout the paper, we use the following conventional mathematical notations. For any matrix  $\mathbf{a} \in \mathbb{R}^{m \times n}$ , we use  $\mathbf{a}_i \in \mathbb{R}^n$  to denote  $i$ -th row vector and  $\mathbf{a}_{:,j} \in \mathbb{R}^m$  to denote the  $j$ -th column vector. The “ $\cdot$ ” is used for inner product of vectors, whereas “ $\odot$ ” is for the Hadamard (entry-wise) product, i.e.,  $\mathbf{a} \odot \mathbf{b} = (a_1b_1, a_2b_2, \dots)$ . For notational convenience, we often write inner product as  $\mathbf{a} \cdot \mathbf{b}$  without explicitly using the transpose notation.

We consider a game with a unit of continuum researchers and finitely many publication venues, denoted by set  $V = \{1, \dots, k\}$ . Any researcher is characterized by a *type*  $\theta \in \mathbb{R}$ , the researcher’s impact factor, determined by her publication history, research taste, productivity, etc. Let  $\Theta \subset \mathbb{R}$  denote the set of all possible researcher types, which is assumed to be discrete. A generic researcher type is denoted as  $\theta_i \in \Theta$  and has mass  $\mu(\theta_i)$ . We sort the parameters  $\theta_1 < \theta_2 < \dots < \theta_n$  increasingly. We assume a researcher’s type  $\theta_i$  does not change over the short term considered in our Publication Choice Problem, an assumption justified by a simulation in which we initialize the venues’ impact at 50 random starting points. In all cases, the simple best-response dynamics converge rapidly—within only a few iterations—to the equilibrium of the game (Table 2). The full experimental setup is described in Appendix A.4.

# Rounds to Convergence	4	5	6	7
Proportion	2%	8%	86%	4%

Table 2: Rounds before converging to equilibrium.

We consider symmetric strategies in the sense that any two researchers of the same type select the same publication strategy. Let vector  $\mathbf{a}_i \in \mathbb{R}^k$  denote type  $\theta_i$ ’s strategy profile, where each entry  $a_{i,j}$  is the number of publications that a type  $\theta_i$  researcher publishes on venue  $j$ . We write the strategy of all researcher types as matrix  $\mathbf{a}$ .

In our model, a venue impact is tied to the types of researchers who publish on the venue. Before reading each paper in detail, the community tends to pay attention to venues which high-impact researchers publish on. Thus, under strategy profile  $\mathbf{a}$ , the venue impact  $v_j$  is modeled as the weighted average type  $v_j = \frac{(\mathbf{a}_{:,j} \odot \boldsymbol{\mu}) \cdot \boldsymbol{\theta}}{(\mathbf{a}_{:,j} \odot \boldsymbol{\mu}) \cdot \mathbf{1}}$  of researchers who publish on venue  $j$ , where researcher types are weighted by their number of publications on venue  $j$ .

We model the cost for researcher  $i$  to publish one paper on venue  $j$  as  $c_{i,j}$ . Venues are sorted by their competitiveness in publication. Naturally, a more competitive venue is assumed to have a higher publication cost for every type of researchers.<sup>2</sup> Thus,  $c_{i,j}$  increases in venue index  $j$  for each researcher of type  $i$ .

<sup>2</sup>We do not model paper acceptance or rejection in our model. When the acceptance/rejection decision is less random, the cost is naturally deterministic. When the acceptance/rejection decision

A Publication Choice Problem hence is specified by a tuple  $(V, \Theta, \mu, \mathbf{c})$ , with meanings summarized below.

$\boldsymbol{\theta} = (\theta_i)_i$	researcher types (impact factors)
$\boldsymbol{\mu} = (\mu_i)_i$	density on each type $i$ in the community
$\mathbf{a} = (a_{i,j})_{i,j}$	type $i$ ’s num of publications on venue $j$
$\mathbf{c} = (c_{i,j})_{i,j}$	type $i$ ’s cost of publication on venue $j$
$\mathbf{v} = (v_j)_j$	average impact of each venue $j$

Table 3: Notations for an Publication Choice Problem.

The researcher derives utility from publishing on venues with higher impacts. Formally, fixing venue impacts  $\mathbf{v}$ , the researcher of type  $\theta_i$  publishes  $\mathbf{a}_i = (a_{i,j})_j$  on each venue  $j$  and gains utility  $u_i(\mathbf{a}_i, \mathbf{v}) = (\mathbf{a}_i^\alpha \cdot \mathbf{v}^\beta)^{\frac{1}{\beta}}$ . Following an axiomatic characterization about how to count research impact due to Perry and Reny (2016), we assume the utility function takes the form of a  $\beta$ -norm. Specifically, Perry and Reny (2016) proves that if a researcher’s impact function over publications satisfies monotonicity, independence, depth relevance, and scale invariance, then the only function format is the  $\beta$ -norm on the vector of citation numbers. Parameters  $\alpha, \beta$  decide the marginal utility gained from a publication and are homogenous across researchers. In the utility function,  $\mathbf{a}_i^\alpha$  captures how a researcher normalizes and counts publications, while  $\mathbf{v}^\beta$  calculates the utility gained from each publication count. We make the following natural assumptions on parameters  $\alpha, \beta$ :

- $\alpha \in (0, 1)$ , meaning that the researcher normalizes publication counts on the same venue in a marginally decreasing way.
- $\beta > 1$ , meaning that the utility  $\mathbf{v}^\beta$  from one publication count is marginally increasing in the venue impact. The reason for this is that a researcher is known for and typically cares more about her most impactful works.

**Best Responses.** Upon observing venue impacts  $\mathbf{v}$ , researchers strategize to maximize the utility from publication. The publication game in real life evolves dynamically: the researchers observe the venue impacts from the previous round and decide their publication strategies for the current round; the venue impacts are updated from researchers’ strategies and are observed again; a new loop of researcher best response starts. We define the best response for each researcher: fixing the venue impacts, each researcher  $i$  optimizes her publication utilities, subject to a total cost budget constraint, in Program (1). We normalize the time budget of all researchers to 1. Note that this normalization also works for PIs with a large research lab, since a larger lab does not

shows a higher degree of inconsistency and arbitrariness, such as in NeurIPS (Cortes and Lawrence 2021), the cost captures the expected cost for one publication. As long as the researcher pays enough cost to surpass the venue’s basic quality requirement, she can exploit the randomness in acceptance by submitting the same paper to similar venues until the paper gets accepted.

increase the PI's budget but lowers publication cost.<sup>3</sup>

$$\max_{\mathbf{a}_i} \mathbf{a}_i^\alpha \cdot \mathbf{v}^\beta \quad \text{s.t. } \mathbf{a}_i \cdot \mathbf{c}_i \leq 1 \quad (1)$$

**The Equilibrium.** We study the following natural equilibrium concept with a continuum of researchers, adapted from (Mas-Colell and Vives 1993). With a continuum of researchers, the venue impacts are unaffected by any single researcher's strategy. The equilibrium condition requires that when all researchers are best responding, their strategies are consistent with the venue impacts. An equilibrium can be viewed as a fixed point in the dynamic loop consisting of best responses.

**Definition 1** (Equilibrium with a Continuum of Researchers). *For any Publication Choice Problem  $(V, \Theta, \mu, \mathbf{c})$ , an action profile  $\mathbf{a} = \{a_{i,j}\}_{i,j}$  and venue impacts  $\mathbf{v}$  are a pure-strategy equilibrium if they satisfy the following two conditions:*

- **[Best Response]** *The strategy  $\mathbf{a}_i^*$  of any type- $\theta_i$  researcher is a best response, i.e. solves Program (1).*
- **[Consistency of Venue Impact]** *The venue impacts are consistent with researchers' strategies:*

$$v_j = \frac{(\mathbf{a}_{:,j}^* \odot \boldsymbol{\mu}) \cdot \boldsymbol{\theta}}{(\mathbf{a}_{:,j}^* \odot \boldsymbol{\mu}) \cdot \mathbf{1}}, \quad \text{for any venue } j \quad (2)$$

### 3 Properties of Publication Choice Problems

We analyze the equilibrium properties of the Publication Choice Problem. In Section 3.1, we derive the closed form of the best response problem for each researcher. In Section 3.3, we derive properties of Publication Choice Problem at equilibrium.

#### 3.1 Characterizing Researchers' Best Responses

We start by characterizing a researcher's best response to the venue impacts as a solution to Program 1. As shown in Lemma 3.1, in the best response strategy of each researcher, the published amount of papers is proportional to 1) the utility  $v_j^{\frac{\beta}{1-\alpha}}$  of publication count on each venue, and 2) the marginal cost  $(c_{i,j})^{\frac{1}{\alpha-1}}$  of one normalized publication count  $a_{i,j}^\alpha$  in utility.

**Lemma 3.1.** *Let  $\mathbf{v}$  be the venue impact vector. Then the best response of any researcher of type  $\theta_i$  can be characterized in closed-form as follows:*

$$a_{i,j} = \frac{(c_{i,j})^{\frac{1}{\alpha-1}} \cdot v_j^{\frac{\beta}{1-\alpha}}}{c_i^{\frac{\alpha}{\alpha-1}} \cdot v^{\frac{\beta}{1-\alpha}}}$$

The proof of Lemma 3.1 derives the closed form solution for Program 1 and is deferred to Appendix B.1.

<sup>3</sup>Program (1) admits the same optimal solution as when the objective is the utility  $(\mathbf{a}_i^\alpha \cdot \mathbf{v}^\beta)^{\frac{1}{\beta}}$ .

#### 3.2 Natural Properties and Model Sanity Check

As a warm-up, this subsection exhibits a few natural properties of the equilibrium hence also serves as a sanity check for the validity of our model above before we dive into more evolved analysis afterwards.

Before analyzing the equilibrium of the Publication Choice Problem, we introduce the following key assumption, the *Monotone Cost Ratio* (MCR). Specifically, we assume the relative cost between a low type and a high type increases at higher/better venues.

**Assumption 1** (Monotone Cost Ratio (MCR)). *The cost ratio between a low type and a high type increases with the venue index  $j$ , i.e. for all types  $\theta_i < \theta_{i'}$  of researcher, and all venues  $j < j'$ ,  $\frac{c_{i,j}}{c_{i',j}} < \frac{c_{i,j'}}{c_{i',j'}}$ .*

Intuitively, the MCR assumption means that high-type researchers gain more *relative* advantage on more selective publication venues. This is a widely adopted assumption in principal-agent problem (Jehle and Reny 2011) and mechanism design (Bergemann and Välimäki 2002) to describe the advantage of a more skilled agent. In these economic applications, MCR is assumed to capture the intuition that a higher or more qualified agent type will have increasing advantage in generating higher-quality outcomes/products (sometimes also known as *Monotone Likelihood Ratio*, or MLR, when exerted cost leads to ordered probabilistic outcomes). Notably, MLR is a widely adopted assumption in economic applications (see, e.g., the textbook (Jehle and Reny 2011)) and even motivated much statistical study on testing MLR properties (Karlin and Rubin 1956).

The following Proposition 3.2 shows that the MCR property may be an intrinsic reason that different publication venues often end up having different average impact in reality — if the ratio  $\frac{c_{i,j}}{c_{i',j}}$  was the same on different venue  $j$ , then all venues will have the same average impact at *any* equilibrium of the game. The proof of Proposition 3.2 is deferred to Appendix B.2.

**Proposition 3.2.** *Suppose the cost ratio  $\frac{c_{i,j}}{c_{i',j}} = c(i; i')$  is a constant that is independent of the venue  $v_j$ . Then the Publication Choice Problem admits a unique equilibrium in which all venues have the same average impact.*

Before proceeding to the main equilibrium analysis, we highlight several natural properties that a reasonable model of publication choice should satisfy. These results serve as a sanity check for our modeling assumptions; formal statements and proofs are deferred to Appendix B:

- **Monotone venue impacts** Under Assumption 1, more competitive venues receive higher impact when all researchers best respond to the observed venue impacts (See Proposition B.2 in Appendix B.3).
- **Scale invariance of costs** Scaling the entire cost matrix  $\mathbf{c}$  by the same positive factor leaves the equilibrium impact unchanged (See Observation B.5 in Appendix B.4).
- **Asymmetric growth of types** In the binary-type setting, increasing the density of high-type researchers raises the equilibrium impact of all venues; increasing the density of low-type researchers lowers it (See Appendix B.7).

### 3.3 Equilibrium Existence and Their Properties

We now turn to prove the general existence of an equilibrium and its various properties. Moreover, we show that the number of publications on the top venue is monotone in researcher impact type. However, the total number of publications across all venues is not necessarily monotone — that is, there exist simple instances where a researcher with lower impact has a larger number of publications in total.

We start our study by showing the existence of equilibria. The proof of Proposition 3.3 is deferred to Appendix B.5.

**Proposition 3.3.** *Every Publication Choice Problem admits a pure-strategy equilibrium.*

Our main finding of this subsection is the following property which shows that the number of publications in the most competitive venue is indeed monotonically increasing in a researcher’s type and, moreover, this property holds in a much stronger sense than merely at equilibrium. That is, it holds for any  $\mathbf{a}_i^* = \arg \max_{\mathbf{a}_i \geq 0: \mathbf{a}_i \cdot \mathbf{c}_i \leq 1} (\mathbf{a}_i)^\alpha \cdot \mathbf{v}^\beta$  that is a best response to some venue impact vector  $\mathbf{v}$ , regardless  $\mathbf{v}$  is an equilibrium vector or not. This in some sense justifies why some of the popular ranking systems (e.g., the *csranks.org*) choose to rank institutes’ impact based a selected set of top venues. Recall that Proposition B.2 says the most competitive venue coincides with the venue with the highest impact. The proof of Theorem 3.1 is deferred to Appendix B.9.

**Theorem 3.1.** *Under Assumption 1, consider any venue impacts  $\mathbf{v}$  and let  $\mathbf{a}_i^*$  be researcher type  $i$ ’ best responses to  $\mathbf{v}$ . Then  $\alpha_{i,k}^* > \alpha_{i',k}^*$  for any two researcher types  $i$  and  $i'$  with  $\theta_i > \theta_{i'}$  ( $k$  is the most competitive venue).*

**Is the number of publications monotone in impact?** In the academic realm, publication count is commonly used as one of the proxies for researchers’ impact.<sup>4</sup> This prompts an examination of the extent to which the quantity of publications is related to a researcher’s type. The following example show that the total number of publications is generally not monotone in a researcher’s impact.

**Proposition 3.4.** *There exists an Publication Choice Problem such that the total number of publications is not weakly increasing in one’s type under the equilibrium.*

*Proof Sketch.* We construct an example with two researcher types and two venues. The high type allocates more effort to the more costly competitive venue, resulting in a lower total number of publications. The detailed parameterization of this example is provided in Appendix B.8.  $\square$

## 4 The Binary-Type Case: Equilibrium Uniqueness and Spotlight Effects

In this section, we turn to a fundamental special case of the binary researcher types, i.e., a high type and a low type.

<sup>4</sup>Whether this is a right metric is out of the scope of this paper’s research, though we do observe the increasing use of the *csranks.org* website, which counts the number of publications at (only) a selected set of top venues, as a proxy for different institutes’ impact in different fields.

It turns out that in this case, we are able to further show the uniqueness of its equilibrium. This uniqueness enables clearer analysis of equilibrium properties and the effects of introducing “spotlight” acceptance, which has become increasingly popular in today’s AI/ML community. Our theoretical results further show how one venue switching to spotlight labeling may lead to unintended consequences on the impact of other venues.

### 4.1 Equilibrium Uniqueness

We show the pure-strategy equilibrium is unique when there is a non-competitive venue that randomly or uniformly accepts all papers. For example, all researchers have the option to publish their paper drafts permanently on Arxiv and not on any other conferences or journals.

**Assumption 2** (Non-competitive Venue). *The least-competitive venue 1 is non-competitive. That is, the cost for publication is the same for all types:  $c_{i,1} = c_1, \forall i$ .*

Before proving the uniqueness of equilibrium, we need to introduce the characteristic function of the Publication Choice Problem. Let  $x = \frac{\alpha_{H,1}}{\alpha_{L,1}}$  and  $\tilde{\mu} = \frac{\mu_H}{\mu_L}$ . We normalize the impact for the low type to be 1, and the impact for the high type  $\theta$ . We define the following characteristic function.

**Definition 2** (Characteristic Function). *Given an Publication Choice Problem, the characteristic function of the choice problem is defined by*

$$f(x) = \sum_l (c_{H,l})^{\frac{\alpha}{\alpha-1}} \cdot v_l^{\frac{\beta}{1-\alpha}} \left( x - \frac{c_{L,1}}{c_{H,1}} \cdot b_l^{-\alpha} \right),$$

$$\text{where each } v_j(x) = \frac{1+b_j x \theta \tilde{\mu}}{1+b_j x \tilde{\mu}} \text{ with } b_j = \left( \frac{c_{H,1} c_{L,j}}{c_{L,1} c_{H,j}} \right)^{\frac{1}{1-\alpha}}.$$

**Idea of the Characteristic Function** Our characteristic function describes the dynamics in the game after researchers best respond. The venue impacts can be characterized by the action ratio on any one venue when all types best respond, where we take the ratio on the first venue. The input to the characteristic function is venue 1’s current action ratio, and the output is the change in the action ratio after all researchers best respond. The sign of the function characterizes the direction of change in all venues’ impact. The zero point corresponds to an equilibrium, where the action ratios stop updating after best response.

In the following lemma, we summarize four key properties of the characteristic function, and we defer their formal proofs to Appendix B.6.

**Lemma 4.1.** *[Properties of the Characteristic Function] The following four properties about the characteristic function hold.*

1. *A binary-type Publication Choice Problem is in equilibrium if and only if  $f\left(\frac{\alpha_{H,1}}{\alpha_{L,1}}\right) = 0$ .*
2. *If  $f\left(\frac{\alpha_{H,1}}{\alpha_{L,1}}\right) < 0$ , when researchers update their actions in response to current venue impacts, the impact of all venues will increase after update;*
3. *If  $f\left(\frac{\alpha_{H,1}}{\alpha_{L,1}}\right) > 0$ , when researchers update their actions in response to current venue impacts, the impact of all venues will decrease after update.*

4. Under assumption 1 and assumption 2, the characteristic function  $f$  is convex in  $x$ .

With Lemma 4.1, we can prove the uniqueness of the equilibrium by showing  $f = 0$  admits a unique solution. The full proof is deferred to Appendix B.12.

**Theorem 4.1.** *Any binary-type Publication Choice Problem with a non-competitive venue (Assumption 2) admits a unique pure-strategy equilibrium.*

We conjecture the equilibrium is also unique under many-type settings (see Appendix A.2 for empirical evidence). We leave its formal proof as future work.

**Conjecture 4.2** (Uniqueness of equilibrium under many-type setting). *We hypothesize that the pure-strategy equilibrium is unique when there is a non-competitive venue that randomly or uniformly accepts all papers.*

Equilibrium uniqueness benefits comparative statics analysis and enables more interpretable policy insights in Section 4.2.

## 4.2 The Effect of Discriminative Acceptance via Spotlight Labeling

In this section, we first introduce a variant of Publication Choice Problem for the impact of papers with spotlight labeling. Many publication venues nowadays started selectively labeling publications as “spotlight” publications. Selected spotlight papers are often tagged in poster sessions, or presented in an oral session in addition to the poster session. The spotlight labeling attracts more attention from the research community, leading to a higher impact gained on selected papers. We analyze the effect of switching to spotlight labeling on the research community. Our analysis of the equilibrium is restricted to a binary-type setting where we can prove the uniqueness of equilibrium for the same reason as in previous sections. We are able to compare the effect of spotlight labeling in equilibrium only when the equilibrium uniquely exists.

Unlike the establishment of a new venue of higher impact, the impact attributed to a spotlight publication is intrinsically linked to the venue’s existing impact. Suppose the program committee of venue  $j$  decides to separately label some papers as “spotlight”, with spotlight papers a  $\frac{1}{\Omega_j}$  of the regular papers ( $\Omega_j > 1$ ). Our model assumes that the impact of spotlight papers is decided by the average impact  $v_j$  on the regular session and fraction  $\frac{1}{\Omega_j}$  of spotlight papers because the majority of the audience are regular session authors. On average, each spotlight paper gains an impact of  $\gamma(\Omega_j)$  times the impact of a paper in the regular session, where the labeling effect  $\gamma(\Omega_j) > 1$  is determined by  $\Omega_j$ . We use  $\gamma(\Omega)$  to denote the vector  $(\gamma(\Omega_1), \dots, \gamma(\Omega_k))$ . In Appendix B.10, we use empirical citation data on CVPR to justify our assumption that  $\gamma(\Omega_j) > 1$ . Since the spotlight impact is uniquely pinned down by the impact on regular session  $v_j$  and spotlight ratio  $\Omega_j$ , we can express the overall venue equilibrium impact as  $\frac{\Omega_j + \gamma(\Omega_j)}{1 + \Omega_j} v_j$ .

We characterize the properties of the equilibrium when a venue switches to spotlight labeling. The design space

for the venue organizer is the cost  $c_{i,j}^S, \forall i$  and the fraction  $1/\Omega_j$  of spotlight publication (by changing paper selection rules). Let  $\mathbf{a}_{:,j}^S = (a_{i,j}^S)_j$  be the vector of spotlight publications by all agents on venue  $j$ . While the average impact of the spotlight papers is  $\gamma(\Omega_j)v_j$ , the spotlight papers have an actual average impact without spotlight labeling of  $v_j^S = \frac{(\mathbf{a}_{:,j}^S \odot \boldsymbol{\mu}) \cdot \boldsymbol{\theta}}{(\mathbf{a}_{:,j}^S \odot \boldsymbol{\mu}) \cdot \mathbf{1}}$ . When choosing from the design space, the organizer faces constraints, including the following.

- The actual impacts of the spotlight papers are higher than regular papers. The actual impacts are the average impact of a hypothetical venue, assuming spotlight papers are selected for this separate hypothetical venue instead of spotlights of the existing venue.

$$\frac{(\mathbf{a}_{:,j}^S \odot \boldsymbol{\mu}) \cdot \boldsymbol{\theta}}{(\mathbf{a}_{:,j}^S \odot \boldsymbol{\mu}) \cdot \mathbf{1}} > \frac{(\mathbf{a}_{:,j} \odot \boldsymbol{\mu}) \cdot \boldsymbol{\theta}}{(\mathbf{a}_{:,j} \odot \boldsymbol{\mu}) \cdot \mathbf{1}}; \quad (3)$$

- It is harder to publish a paper labeled “spotlight” for all types, i.e.  $c_{:,j}^S \geq c_{:,j}$ .

	Venue $j$	Spotlight $j$	
Research Impact	$v_j$	Spotlight Impact (after labeling) $\gamma(\Omega_j) \cdot v_j$	Actual Impact $\frac{(\mathbf{a}_{:,j}^S \odot \boldsymbol{\mu}) \cdot \boldsymbol{\theta}}{(\mathbf{a}_{:,j}^S \odot \boldsymbol{\mu}) \cdot \mathbf{1}}$
Action	$\mathbf{a}_{:,j}$	$\mathbf{a}_{:,j}^S$	
Cost	$\mathbf{c}_{:,j}$	$\mathbf{c}_{:,j}^S$	

Table 4: Notations for a venue  $j$  with spotlight labeling.

We make the following similar monotone cost ratio assumption on the spotlight session due to a similar reason to Assumption 1. In Appendix B.11, we show if publishing a spotlight paper is relatively the same hard as publishing a regular paper, the actual impact of spotlight papers as a new venue will be the same as regular papers. This violates Constraint 3 that spotlight papers should gain more actual research impact than regular papers on average. Assumption 3 states that publishing a spotlight paper should be relatively harder for lower types than a regular paper.

**Assumption 3** (Monotone Spotlight Cost Ratio). *The relative cost for any low type to publish a spotlight paper is higher than the relative cost to publish a regular paper. i.e.*

$$\text{for any two types } \theta_i < \theta_{i'}, \frac{c_{i,j}^S}{c_{i',j}^S} > \frac{c_{i,j}}{c_{i',j}}.$$

Intuitively, the expected cost of publishing a spotlight paper can dynamically change with the number of publications. In Appendix B.13, we solve the researcher’s utility maximization problem with spotlight session and show that the cost of publishing a spotlight paper can be fixed once the cost of publishing a regular paper is fixed.

### Characterization of equilibrium with spotlight labeling

In Section 4.2, we provide characterizations of the equilibrium after switching to spotlight labeling. We focus on a binary-type Publication Choice Problem in this section. Corollary 4.2 shows the equilibrium is unique. The proof of Corollary 4.2 is deferred to Appendix B.14.

**Corollary 4.2.** Consider any binary-type Publication Choice Problem with one venue using the spotlight labeling. Under Assumption 2 and Assumption 3, there exists a unique pure-strategy equilibrium.

The following theorem shows that, if the organizer cares about absolute research impact but not relative impact in the community, then less competitive venues are better off switching to spotlight labeling. Otherwise, if the venue’s regular venue is competitive enough, the spotlight session attracts too much research impact that it hurts the average research impact on every venue in the community.

**Theorem 4.3.** Consider a binary-type Publication Choice Problem under Assumption 2 and Assumption 3. Then there exists a threshold venue  $j_0$  such that

- if a venue  $j \geq j_0$  (more competitive) switches to spotlight labeling, the equilibrium impact of all venues decrease;
- if a venue  $j < j_0$  (less competitive) switches to spotlight labeling, the equilibrium impact of all venues increase.

The proof of Theorem 4.3 reduces the problem to the scaling effect of the equilibrium. The spotlight publications can be viewed as attracted to a separate venue, which changes the ratio of remaining types in the community. When the remaining types in the community scale with different proportion, the equilibrium impact of venues shifts monotonically. We defer the proof of Theorem 4.3 to Appendix B.15.




Figure 1: Impact factors of regular sessions with and without a venue switching to spotlight labeling. Higher venue indices correspond to more competitive venues.

**Empirical: threshold effect of spotlight labeling holds for many-type settings.** Under Conjecture 4.2, we empirically show the threshold effect in a simulation with five researcher types and three venues; the full setup appears in Appendix A.3. Figure 1 compares the baseline equilibrium impacts (solid blue line) with those obtained when a single venue adopts spotlight labeling (dashed lines). When a more competitive venue introduces spotlight labeling—shown by the red (venue 3) and green (venue 2) dashed lines—the impact factors at all regular sessions fall below the baseline. This occurs because the spotlight session of a competitive venue attracts disproportionately many high-type researchers away from regular sessions. In contrast, when a less competitive venue adopts spotlight labeling—shown by the orange dashed line (venue 1)—the impact at venue 2 rises above the baseline. These results empirically generalize the threshold effect to the many-type setting.

Recall that Constraint 3 requires that spotlight papers have a higher impact on average. One obvious strategy that organizers may use is to only select papers by researchers with

high research impact into the spotlight session, i.e. setting costs  $c_{H,j,S} < \infty$  and  $c_{L,j,S} = \infty$ . We note that this strategy leads to high-type researchers less willing to publish on regular venues, and decreases the impact of all regular venues in the community.

**Corollary 4.3.** Under Assumption 2 and Assumption 3, for a binary-type Publication Choice Problem with  $k$  venues, where some venue  $j$  switches to spotlight labeling, If venue  $j$  only labels papers from high-impact type  $\theta_H$  researchers as “spotlight”, the equilibrium research impacts of all regular venues will decrease.

The proof follows directly from the proof of Theorem 4.3. Corollary 4.3 implies that the organizer should diversify the set of authors with “spotlight” papers.

## 5 Simulations, Conclusions, and Future Work

**Summary of Simulation Results** We empirically study the Publication Choice Problem to validate our modeling and results. Appendix A.1 describes simulation setups. Appendix A.2 checks the uniqueness of equilibrium under many-type setting. Appendix A.3 and Appendix A.4 describe the experimental setups for the many-type threshold effect and the fast convergence result, respectively. Finally, Appendix A.5 and Appendix A.6 analyze how relative costs and spotlight ratios influence the equilibrium outcomes

**Conclusion and Future Work** This paper proposes a game-theoretic model, the Publication Choice Problem, that explains the interplay between researchers’ publication choice and the evolution of publication venues’ impact. We study the properties of the game in equilibrium from an observer’s perspective of the research community.

Our results can be divided into two sets: results that an observer of the research community should expect from a game-theoretic model: the equilibrium existence and the scaling effect, which justify our model choice; results that shed light on the publication patterns: the monotonicity of publication number in researcher type and the effect of spotlight labeling. In future work, we will consider the optimization problem of the venue organizers.

We outline several future directions both theoretically and empirically. On the theory side, we leave the formal proof for equilibrium uniqueness (Conjecture 4.2) and the threshold effect (Theorem 4.3) under many-type setting as future work. On the empirical side, our model is not limited to analyzing the AI publication market. Observation B.5 in Appendix suggests that understanding the relative cost of publication in different fields allows our model to predict equilibrium outcomes in other disciplines. One future work is to gather more data and estimate the spotlight advertisement effect  $\gamma$ . A closed form function  $\gamma$  may lead to new theoretical conclusions and practical insights. Given the generality of our model, it could assist a wide range of academic communities in optimizing the impact of their publication venues.

## Acknowledgements

Haifeng Xu acknowledges support from the AI2050 program at Schmidt Sciences (Grant G-24-66104) and the Army Research Office Award W911NF-23-1-0030. This work was done when Haichuan Wang was an undergraduate student at University of Chicago and when Yifan Wu was a PhD student at Northwestern University, supported by NSF ECCS 2216970, under the IDEAL Summer Research Exchange Program.

## References

- Aumann, R. J. 1975. Values of markets with a continuum of traders. *Econometrica: Journal of the Econometric Society*, 611–646.
- Aziz, H.; Micha, E.; and Shah, N. 2023. Group fairness in peer review. *Advances in neural information processing systems*, 36: 64885–64895.
- Bergemann, D.; and Välimäki, J. 2002. Information acquisition and efficient mechanism design. *Econometrica*, 70(3): 1007–1033.
- Blonski, M. 1999. Anonymous games with binary actions. *Games and Economic Behavior*, 28(2): 171–180.
- Boehmer, N.; Bredereck, R.; and Nichterlein, A. 2022. Combating collusion rings is hard but possible. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 4843–4850.
- Cortes, C.; and Lawrence, N. D. 2021. Inconsistency in conference peer review: revisiting the 2014 neurips experiment. *arXiv preprint arXiv:2109.09774*.
- Ductor, L.; Goyal, S.; van der Leij, M. J.; and Paez, G. N. 2020. On the Influence of Top Journals. Technical report, Tinbergen Institute Discussion Paper.
- Frank, M. R.; Wang, D.; Cebrian, M.; and Rahwan, I. 2019. The evolution of citation graphs in artificial intelligence research. *Nature Machine Intelligence*, 1(2): 79–85.
- Friedman, E. J. 1996. Dynamics and rationality in ordered externality games. *Games and Economic Behavior*, 16(1): 65–76.
- Fromm, M.; Faerman, E.; Berrendorf, M.; Bhargava, S.; Qi, R.; Zhang, Y.; Dennert, L.; Selle, S.; Mao, Y.; and Seidl, T. 2021. Argument mining driven analysis of peer-reviews. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4758–4766.
- Heckman, J. J.; and Moktan, S. 2020. Publishing and promotion in economics: The tyranny of the top five. *Journal of Economic Literature*, 58(2): 419–470.
- Jecmen, S.; Zhang, H.; Liu, R.; Shah, N.; Conitzer, V.; and Fang, F. 2020. Mitigating manipulation in peer review via randomized reviewer assignments. *Advances in Neural Information Processing Systems*, 33: 12533–12545.
- Jehle, G.; and Reny, P. 2011. *Advanced Microeconomic Theory*. Harlow: Pearson, 3 edition.
- Karlin, S.; and Rubin, H. 1956. The theory of decision procedures for distributions with monotone likelihood ratio. *The Annals of Mathematical Statistics*, 272–299.
- Lasry, J.-M.; and Lions, P.-L. 2007. Mean field games. *Japanese journal of mathematics*, 2(1): 229–260.
- Lauriere, M.; Perrin, S.; Girgin, S.; Muller, P.; Jain, A.; Cabannes, T.; Piliouras, G.; Pérolat, J.; Elie, R.; Pietquin, O.; et al. 2022. Scalable deep reinforcement learning algorithms for mean field games. In *International conference on machine learning*, 12078–12095. PMLR.
- Lipton, Z. C.; and Steinhardt, J. 2019. Troubling Trends in Machine Learning Scholarship: Some ML papers suffer from flaws that could mislead the public and stymie future research. *Queue*, 17(1): 45–77.
- Manzoor, E.; and Shah, N. B. 2021. Uncovering latent biases in text: Method and application to peer review. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4767–4775.
- Mas-Colell, A.; and Vives, X. 1993. Implementation in Economies with a Continuum of Agents. *The Review of Economic Studies*, 60(3): 613–629.
- Meir, R.; Lang, J.; Lesca, J.; Mattei, N.; and Kaminsky, N. 2021. A market-inspired bidding scheme for peer review paper assignment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4776–4784.
- Milchtaich, I. 1996. Congestion models of competition. *The American Naturalist*, 147(5): 760–783.
- Payan, J.; and Zick, Y. 2022. I Will Have Order! Optimizing Orders for Fair Reviewer Assignment. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 1711–1713.
- Perrin, S.; Pérolat, J.; Laurière, M.; Geist, M.; Elie, R.; and Pietquin, O. 2020. Fictitious play for mean field games: Continuous time analysis and applications. *Advances in neural information processing systems*, 33: 13199–13213.
- Perry, M.; and Reny, P. J. 2016. How to count citations if you must. *American Economic Review*, 106(9): 2722–2741.
- Roughgarden, T.; and Tardos, É. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and economic behavior*, 47(2): 389–403.
- Shah, N. B. 2021. Systemic challenges and solutions on bias and unfairness in peer review. *Preprint [http://www.cs.cmu.edu/nihars/preprints/Shah\\_Survey\\_PeerReview.pdf](http://www.cs.cmu.edu/nihars/preprints/Shah_Survey_PeerReview.pdf)*, 2.
- Spence, M. 1978. Job market signaling. In *Uncertainty in economics*, 281–306. Elsevier.
- Stelmakh, I.; Shah, N.; and Singh, A. 2021a. PeerReview4All: Fair and accurate reviewer assignment in peer review. *Journal of Machine Learning Research*, 22(163): 1–66.
- Stelmakh, I.; Shah, N. B.; and Singh, A. 2021b. Catch me if i can: Detecting strategic behaviour in peer assessment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4794–4802.
- Su, W. J. 2021. A Truthful Owner-Assisted Scoring Mechanism. *Neural Information Processing Systems (NeurIPS)*, 28: 28.
- Wang, D.; Song, C.; and Barabási, A.-L. 2013. Quantifying long-term scientific impact. *Science*, 342(6154): 127–132.

Wright, J. R.; Thornton, C.; and Leyton-Brown, K. 2015. Mechanical TA: Partially automated high-stakes peer grading. In *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, 96–101.

Wu, J.; Xu, H.; Guo, Y.; and Su, W. 2024. A Truth Serum for Eliciting Self-Evaluations in Scientific Reviews. *arXiv preprint arXiv:2306.11154*.

Zhang, Y.; Schoenebeck, G.; and Su, W. 2024. Eliciting Honest Information From Authors Using Sequential Review. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9977–9984.

Zhang, Y.; Yu, F.-Y.; Schoenebeck, G.; and Kempe, D. 2022. A System-Level Analysis of Conference Peer Review. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 1041–1080.